Published at 2005 International Symposium of Robotics Research, October 2005, San Francisco, California, USA

Tradeoffs Between Directed and Autonomous Driving on the Mars Exploration Rovers

Jeffrey J. Biesiadecki and Chris Leger and Mark W. Maimone
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA USA
first.last@jpl.nasa.gov

October 5, 2005

Abstract

NASA's Mars Exploration Rovers (MER) have collected a great diversity of geological science results, thanks in large part to their surface mobility capabilities. The six wheel rocker/bogie suspension provides driving capabilities in many distinct terrain types, the onboard IMU measures actual rover attitude changes (roll, pitch and yaw, but not position) quickly and accurately, and stereo camera pairs provide accurate position knowledge and/or terrain assessment. Solar panels generally provide enough power to drive the vehicle for at most four hours each day, but drive time is often restricted by other planned activities. Driving along slopes in nonhomogeneous terrain injects unpredictable amounts of slip into each drive. These restrictions led us to create driving strategies that maximize drive speed and distance, at the cost of increased complexity in the sequences of commands built by human Rover Planners each day.

The MER rovers have driven more than a combined 10 kilometers over Martian terrain during their first 21 months of operation using these basic modes. In this paper we describe the strategies adopted for selecting between human-planned directed drives versus rover-adaptive Autonomous Navigation and Visual Odometry drives.

1 Background

NASA successfully landed two mobile robot geologists on the surface of Mars in January 2004: the Spirit and Opportunity Mars Exploration Rovers (MER). Their primary goal was to find evidence of past water at Gusev Crater and Meridiani Planum, two geologically distinct sites on opposite sides of the planet. Each rover was instrumented with a suite of tools for remote sensing (multi-filter and stereo camera pairs and a thermal emission spectrometer) and *in situ* measurement (5 DOF arm for deploying a grinding Rock Abrasion Tool, Microscopic Imager, Alpha Particle X-ray Spectrometer, and Mössbauer Spectrometer). Although the achievement of their successful landings stands out as a technological tour de force, it was their ability to traverse while on the surface of Mars that enabled both rovers to succeed in their primary goals.

The MER rovers are typically commanded once per Martian solar day (or *sol*). A sequence of commands sent in the morning specifies the sol's activities: what images and data to collect, how to position the robotic arm, and where to drive. At the end of each sol, the rovers send back the images and data human operators will use to plan the next sol's activities. The next sol's mobility commands are selected based on what is known – and what is unknown – about the terrain ahead.

1.1 Rover Mobility Commands

The rovers are driven using three primary modes: low-level commands that specify exactly how much to turn each wheel and steering actuator, directed driving primitives for driving along circular arcs (of which straight line driving and turn-in-place are special cases), and autonomous path selection.

Several types of potential vehicles hazards are checked *reactively*, most of them during Real Time Interrupts (RTIs) which occur 8 times per second. Available checks include Tilt/Pitch/Roll, Northerly Tilt, Rocker/Bogie Suspension Angles, Motor Stalls, Limit Cycle (no forward progress), and Resource Contention.

The rovers maintain an estimate of their local position and orientation updated at 8 Hz while driving. Position is first estimated based on wheel odometry, and orientation is estimated using an Inertial Measurement Unit that has 3-axis accelerometers and 3-axis angular rate sensors. In between driving primitives, the rover can use camera-based Visual Odometry (VisOdom) to correct the errors in the initial wheel odometry-based estimate. VisOdom tracks terrain features in NavCam stereo images and uses the tracking information to estimate true vehicle motion during small steps; the rover can only move roughly 60cm, or turn 15 degrees, before successive NavCam images lack enough overlap to reliably estimate motion [2].

Both directed and path selection modes of driving can make use of onboard stereo vision processing and terrain analysis software to determine whether the rover would encounter geometric hazards along its chosen path.

The computing resources required by these different commands vary greatly. Directed driving commands execute the most quickly (achieving speeds up to 124 m/hour), but also have greater risk since the rover can only count wheel rotations to estimate position, and never looks ahead to evaluate the terrain before driving onto it. AutoNav commands detect and avoid geometric hazards, but only achieve driving speeds from 10 m/hour in obstacle-laden terrain up to 36 m/hour in safe terrain, and also rely on the accuracy of the wheel odometry to track obstacles once they leave the field of view of the cameras. VisOdom commands provide accurate position estimates (but not obstacle detection), and require close spacing between images which limits the top speed to 10 m/hour.

1.1.1 Autonomous Terrain Analysis

When information about nearby terrain is unavailable or uncertain, the rover can be commanded to evaluate terrain safety by performing stereo vision and terrain assessment autonomously. This allows the rover to *predictively* locate traverse hazards and avoid them. The procedure is not summarized here; see [4, 1] for details and [9] for the approach that inspired it.

The rock-strewn terrain encountered by Spirit at Gusev Crater corresponded well to the exponential rock distribution models predicted using data from Viking I, II and Pathfinder missions [5]. The body-mounted 120-degree Field of View (FOV) HazCams were designed with this terrain model in mind, and Spirit has performed all of its autonomous terrain assessment using these cameras. However, the terrain encountered by Opportunity at Meridiani Planum is vastly different. Instead of a wide variety of rocks at many scales, much of the terrain consists of very fine-grained materials; so fine, in fact, that no large scale features can be found in the wide FOV HazCam images at 256x256 resolution. Fortunately, the lack of large scale features implies a lack of large "step" obstacles. So, Opportunity was reconfigured to perform terrain assessment with more narrow FOV NavCam images. Rock and fissure obstacles can still be detected, but the limited FOV means less of the terrain around the obstacle will be understood, which reduces its ability to steer around them autonomously.

All MER surface software runs on a 20 MHz RAD6000 computer under the VxWorks operating system. The slow processor speed, and the sharing of a single address space and cache by dozens of tasks, mean Autonomous Navigation (AutoNav) and VisOdom software run slowly.

1.2 Ground-based Terrain Analysis

Ground-based terrain assessment is generally performed using stereo image pairs taken by any of the three types of stereo camera pairs found on MER vehicles. There are two pairs of wide field-of-view (120 degree, 10cm baseline) Hazard Cameras (HazCams) rigidly mounted 53cm above the ground plane on the front and back sides, one pair of medium field-of-view (45 degree, 20cm baseline) Navigation Cameras (NavCams) mounted 152cm above the ground plane on a pan/tilt head, and one pair of narrow field-of-view (18 degree, 28cm baseline) Panoramic Cameras (PanCams) also mounted 152cm above the ground plane on the pan/tilt head. [8] These cameras take up to 1024x1024 12-bit images that provide information about terrain texture throughout their images, and stereo range-derived terrain shape at different scales: around 0.5m - 5m in the HazCams, 2m - 20m in the NavCams, and 4m - 70m in the PanCams.

The amount of directed driving that can be commanded depends on both the terrain itself and on how much information about the terrain is available. Orbital imagery, while crucial for long-range planning, cannot resolve vehicle hazards like 20cm rocks. So after each long drive, images from each appropriate camera pair are requested.

Downlinked stereo image pairs are processed by an automated pipeline that generates derived products including 3D range maps, texture-mapped terrain meshes, and color overlays indicating terrain properties such as slope and elevation [6]. Rover operators use image-based querying tools to measure ranges to terrain features and estimate distances and rock sizes [3]. For example, a "ruler" tool allows the operator to measure the distance between the 3D points corresponding to two pixels in an image or image mosaic, useful for identifying discrete obstacles such as rocks or steps. Terrain meshes give the operator a geometric understanding of the terrain and of spatial relationships between terrain features and the planned path, and allow simulation of drive sequences to predict drive safety and performance [10]. The raw images are also extremely useful in assessing traversability: operators can readily identify very sandy or very rocky areas that present hazards, though new terrain types always carry an element of uncertainty regarding vehicle performance. In some cases, no image cues enable rover operators to predict the performance of a drive; patches of terrain only a few meters apart, with similar surface texture and geometry, can lead to drastically different traction or sinkage. For example, while driving uphill toward a topographic high point named “Larry’s Lookout” on sol 399, Spirit reached 100% slip (i.e. no forward progress) on a 16 degree slope, but only a few meters further had only 20% slip on a 19 degree slope with no discernible difference in appearance.

Humans are very good at terrain analysis for motion planning. In addition to geometric hazards such as rocks or drop-offs, humans can readily identify and classify new terrain types (e.g., sandy versus rocky slopes) on the basis of appearance alone. In contrast, the MER software does not have any appearance-based terrain analysis capabilities, it only detects geometric obstacles. Nevertheless, the most serious and frequent hazards (rocks, steps, and high-center hazards) can be detected by geometric analysis–assuming sufficient range data is available. At longer ranges (over 15m in NavCam images, and over 50m in PanCam images), range data becomes sparse, making it impossible to rely solely on geometric analysis. The rover is better able to assess nearby hazards, but its lack of a global planner (which the human stands in for during manual drives) can cause the rover to get stuck in cul de sacs.

2 Drive Techniques and Templates

Most drive sequences can be classified as either traverses (covering maximum distance) or approaches (driving to a specific position for subsequent *in situ* arm operations). The techniques used for each drive type are determined based on the time allocated for driving, the amount of terrain visible in imagery, known hazards, and level of uncertainty in rover position given the terrain type. Generally, driving on level ground requires a mix of blind and AutoNav driving, and driving on slopes requires using VisOdom to allow the rover to compensate for unpredictable slip.

2.1 Traversing the Plains

We learned during our initial drives in each terrain that driving on level ground typically leads to accurate and predictable mobility performance; e.g., Spirit only accumulated 3% position error over 2 kilometers of driving [7]. Because of the rover’s limited processing power, drives using autonomous hazard avoidance are several times slower than “blind” (manually-directed) drives. These two facts favor long initial blind drives to achieve the longest drives in the least amount of rover execution time. Human operators can easily identify rocks that are large enough to be hazardous to the rover, and can plan complex paths that avoid them. The firm surfaces found on the plains of Gusev crater often allowed for blind drives of up to 70m.

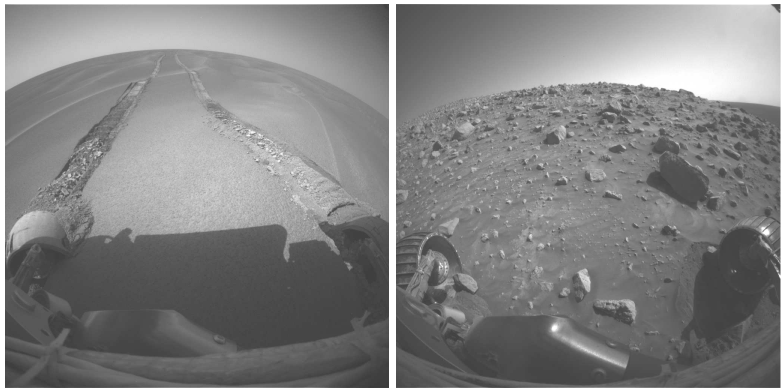

Figure 1: Left: On Sol 446, Opportunity found its wheels more than half buried in sand. Although not a geometric hazard, the ripple of sand on which it stopped kept the human planners busy for weeks. Right: On Sol 454, Spirit terminated its drive early after detecting 90% slip. This image shows rocks that had collected next to the left front wheel.

On the plains of Meridiani, the terrain hazards were quite different and initially allowed for blind drives over 100m. Unlike the Gusev plains, there was a near-total absence of rocks at Meridiani, and until sol 446 (see Figure 1) none of the innumerable sandy ripples posed a threat to the rover. Craters, visible in orbital imagery, and small linear depressions were the most significant hazards for Opportunity. While driving over flat terrain, the rover's suspension does not articulate significantly, which suggested that a measured suspension articulation change could be used to halt driving if the rover were to encounter a depression. In April 2004, the rover's software was upgraded to allow the rover's suspension angles to be checked against preset limits at 8Hz, thus enabling the rover to stop at negative terrain features (i.e., holes) that were not visible a priori. Because the reason for halting a drive (e.g., timeout, suspension check, slip amount, or tilt check) is accessible to the rover sequencing language, a recovery maneuver could be performed whenever the suspension check tripped. The recovery consists of backing up several meters and continuing the drive with AutoNav, since AutoNav is able to detect and avoid negative hazards.

Both rovers use a common strategy at the end of long traverses to acquire necessary images for manipulator operations and turn to a preset heading that minimizes the multi-path interference caused by the rover's mast during communication with Earth or an orbiter. However, this presents a problem for the next sol's IDD operations: since no camera can see the part of the IDD deployment volume under the rover, a front HazCam image pair of the final terrain must be safely acquired 0.5-3m before driving to the rover's final location in order to determine if the IDD can be safely deployed.

The obvious solution is to turn to the desired heading, acquire the image pair, then drive a short distance to the final location. The "guarded arc" drive primitive solves this problem by only executing the post-turn drive segment if the onboard terrain analysis shows that it is safe to do so.

2.2 Driving on Slopes: Mountains and Craters

While most of the distance covered by the rovers was on level ground, most of the sols and most of the approach drives occurred on slopes. The rovers invariably slip when driving on slopes, making VisOdom essential for safe and accurate driving. But using AutoNav along with VisOdom takes roughly twice as much time as VisOdom alone, making the combination impractical for normal use.

This presents a challenge: the rover has the ability to know where it is, but in that mode cannot detect obstacles. Additionally, in steep terrain the rover cannot identify all obstacle classes, since the rover has no means of detecting sandy, high-slip areas in advance. Even otherwise safe areas of moderate slope may represent hazards if there are steeper slopes or rocks downhill, since slippage in moderate slopes could take the rover into dangerous areas. In these cases, the rover operators specify "keep out zones" which will cause the rover to halt driving before a hazard is encountered (e.g., see Figure 3). The rover keeps track of its position using VisOdom (and can preclude driving if VisOdom fails) and can close the loop to correct for slippage. On sol 454, Spirit promptly halted driving after detecting slippage over 90%, and post-drive HazCam images showed several rocks on the verge of falling into the wheels, since the wheels had dug into the terrain by nearly one wheel radius (see Figure 1). The recurrence of high slopes, sandy terrain with intermixed small rocks, and frequent obstacle-sized rocks caused us to retreat and find a new route to the summit of Husband Hill.

2.3 Target Approach

Whereas traverse sequences focus on covering maximum distance over terrain, target approach sequences aim to place the rover at a specific target position and orientation for *in situ* examination of rocks and soil with the rover's manipulator, or less frequently, high-resolution imagery of a distributed or inaccessible target region. The accuracy requirements for positioning the rover for *in situ* work are relatively tight, often within 10cm.

On level ground, directed drive primitives are usually sufficient for accurate target approaches from 2-10m away. On sloped terrain, VisOdom is required to close the loop on the rover's position. After each motion, VisOdom updates the rover's position knowledge, allowing it to correct for slip-induced errors. Conditional sequencing that confirms the current distance to multiple targets is often used in conjunction with visual odometry to accurately approach targets 5-10m away while driving on slopes in the 10 to 20 degree range (e.g., see Figure 2), with the caveat that on surfaces with sufficiently low bearing strength, the rover is mechanically incapable of making direct uphill progress.

3 Relative Merits of Directed/Autonomous Driving

There are significant differences in resource usage between manual and autonomous driving, with execution time and generated data volume being the most obvious. Power is also impacted by execution time, for although the power used by the mobility system is the same whether a trajectory

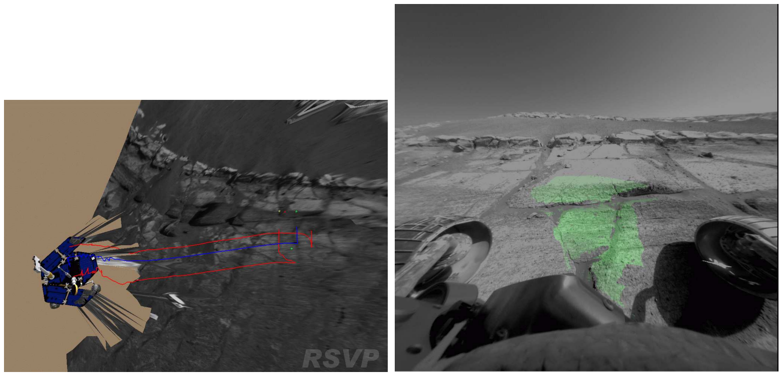

Figure 2: Opportunity's planned 8.7 meter drive along a 20-24 degree slope on Burns Cliff on Sol 304, and the front HazCam view confirming a successful single sol approach. The shaded area shows those parts of the surface reachable by the instrument arm, which includes the light bedrock that was the target of the drive. A combination of VisOdom and conditional sequencing was used to accomplish this drive.

was generated manually or autonomously, the rover's CPU, IMU, and other electronics draw power for the duration of the drive and thus an autonomous drive will require more power than a manual drive of the same distance.

Less obvious differences in resource requirements between manual and autonomous driving also exist. The most significant is planning time: it takes a rover operator more time to identify obstacles and choose appropriate waypoints when sequencing a blind drive than when sequencing a drive using AutoNav (e.g., see Figure 3). During the first few months of the mission, it often took up to 10 hours to build a drive sequence to travel 20-40m across the plains of Gusev. This decreased dramatically later in the mission, often requiring only 2-4 hours to sequence drives over 100m in length on either rover. Still, a directed drive places full responsibility for vehicle safety on the rover operator rather than allowing the rover to safeguard itself, thus requiring more time for manual terrain analysis and waypoint selection. This suggests an obvious trade-off between sequencing time and execution time for directed and autonomous drives.

There is an additional long-term resource trade-off: humans can rapidly adapt their sequences to deal with new terrain types or drive requirements, but changing the onboard software involves a lengthy software development, testing, and uplink process. Instead of a day-to-week turnaround in sequence development, software updates to cope with new terrain and drive techniques occur on a months-to-year cycle.

3.1 Driving into the Unknown

There is one notable circumstance in which the human has no ability to safely select paths: when driving into terrain that has not been imaged. On sol 109, Spirit was commanded to drive over the local horizon 50m distant as it descended from the rim of Missoula Crater. In this case, AutoNav was the

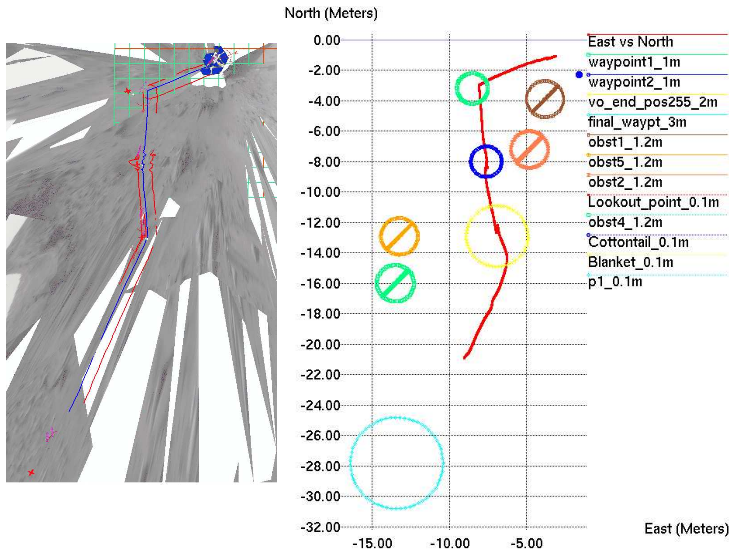

Figure 3: Spirit's Sol 436 drive used a variety of driving modes. A simulation of the planned drive over a 3D terrain mesh is shown on the left, the actual course taken on the right. Circles indicate the waypoints, slashed circles the obstacles and keep-out zones developed by human Rover Drivers by inspecting stereo images and simulating drives over the 3D mesh. Spirit drove south (downward) over 26 actual meters before reaching its time limit. Jagged lines in the course plot above the -12 meter line indicate the discrete jumps resulting from VisOdom updates, those at the -12 meter line show AutoNav backing up to avoid a small ridge blocking its path southwest.

only option available to drive further and use the available time and power, and post-drive images showed AutoNav correctly avoiding large rocks while traversing slopes up to 9 degrees (see Figure 4). Obviously, a high degree of confidence in the hazard avoidance software is needed in situations such as this; *AutoNav has kept both vehicles safe through over 2500 meters of traverse as of August 2005.* Less severe, but more frequent, instances in which humans cannot guarantee rover safety occur when the rover drives beyond the distance at which obstacles can be resolved, or through smaller occluded regions. In practice, even when using AutoNav the rover operator typically chooses waypoints that avoid the most hazardous areas, thus taking advantage of the perceptual strengths of both human and rover.

3.2 Execution

Directed drives have a limited ability to deal with errors or uncertainty in execution. Whereas AutoNav can close the loop on vehicle safety by imaging the terrain that the rover is about to drive through, a directed drive must make the assumption that the rover does not deviate far enough from the planned path to encounter any hazards. For longer drives or in high-slip areas, the rover must be able to deal with accumulated position error, ei-

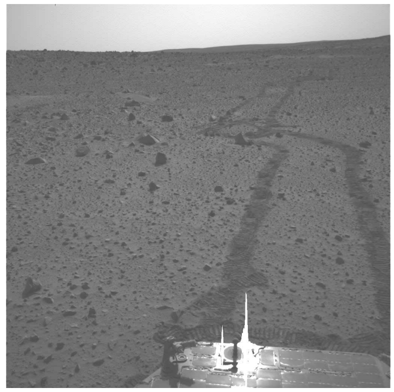

Figure 4: On Sol 109, Spirit avoided obstacles in previously-unseen terrain.

ther through safeguarding itself or by using VisOdom to update its position knowledge. When using VisOdom, the rover operator is responsible for specifying the criteria for halting the drive, since manually sequencing reliable obstacle avoidance is too difficult. Typically, the halting criteria include proximity to known obstacles, the number of times VisOdom has failed to provide a position update, and a threshold on slippage.

Figure 5 summarizes the distance covered and the type of driving modes used for each rover during their first 19 months of operation.

3.3 Adaptation

Mobility performance is uncertain in any novel terrain type and can vary substantially in known terrain types, but humans can quickly learn to steer the rover clear of newly identified hazard types. For example, after Spirit drove through a loose mixture of fine sand and rocks on sol 339, a potato-sized rock jammed in one of the wheels, finally coming out a week later. When the rover encountered similar terrain over 100 sols later, rover operators knew to direct Spirit to check for slippage while driving and stop if the rover became bogged down. Post-drive images after the rover detected over 90% slip showed a similar mixture of sand and rocks, with two rocks having the potential to jam in the wheels, and we subsequently retreated to look for another route (see Figure 1). This sort of perception and adaptation with a single training example is a key strength of manual terrain analysis.

4 Future Work

While Spirit and Opportunity continue to perform well beyond our original expectations, our experience operating the rovers suggests some areas for improvement. Perhaps the most obvious area for improvement is computational efficiency: driving with either VisOdom or AutoNav can slow the rovers' progress by up to an order of magnitude compared to directed drives. Some speedup can likely be obtained by accepting decreased accuracy: one

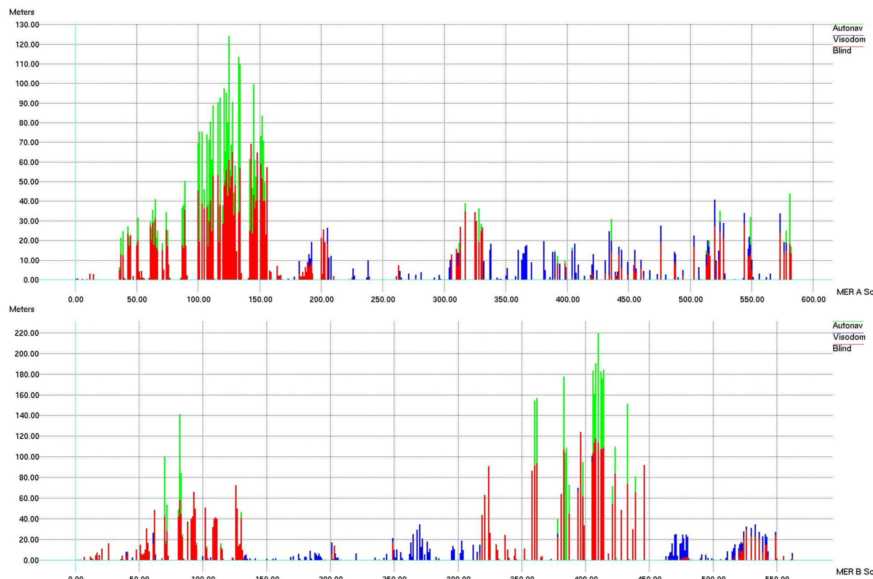

Figure 5: Summary of distances driven by each rover (Spirit above Opportunity) per Sol. AutoNav drives (in green) include any mode in which terrain assessment was done onboard (i.e., both AutoNav and Guarded motion), VisOdom drives (in blue) include both directed and adaptive driving modes but not AutoNav, and Blind drives (in red) include both directed arcs and rover-adapted drives that compensated for yaw changes measured during the drive. The changing quality of the drive types suggests how human and rover driving strategies alike had to adapt to new terrains many times over the course of each mission.

use of VisOdom is to simply detect when the rover is slipping substantially, in which case a precise motion estimate is not required.

Another promising avenue for future work is terrain classification. Our current hazard avoidance software detects only geometric hazards, but areas with weak soil–particularly wind-driven drifts–have proven treacherous for both rovers. The ability to learn what high-slip terrain looks like so that it can be autonomously avoided (even dynamically updating the onboard interpretation of the terrain) would be a great benefit. One potentially useful observation is that slippage is almost always correlated with sinkage, and sinkage can be measured by observing either the wheels or their tracks.

In terms of mobility system development, one area that seems to be underemphasized is precision mobility in natural terrain. For the types of investigation undertaken by Spirit and Opportunity, mere mobility–the ability to traverse a certain-sized obstacle, travel at a certain rate, or climb a certain slope–is not sufficient. The ability to reliably navigate the rover to within centimeters of a desired location, on slopes, near obstacles, and with external constraints on final vehicle heading, has been of the utmost importance in uncovering the water history of Mars.

Flexibility in the rovers' command language and onboard software has been critical in allowing us to encode our ever-changing understanding of the terrain and vehicle performance. While not a traditional robotics problem, it would be beneficial to introduce methods for easily formalizing and re-using new sequence idioms to reduce human errors and speed the sequence design, simulation and validation processes. Writing a sequence is writing a program, and perhaps techniques could be applied from extreme programming and other software development paradigms.

MER software development continues today. Several technologies are being evaluated for possible uplink in mid-2006. These include autonomous *in situ* instrument placement following a successful drive (aka Go and Touch), global path planning to enable intelligent backtracking, visual servoing, and autonomous detection of dust devils and clouds in onboard imagery.

Future vehicles will have faster processors, allowing more advanced terrain analysis and path selection to be performed. But path planning can only be as good as the underlying obstacle avoidance methodology, and if rovers are to become substantially autonomous then appearance-based adaptive terrain analysis will also be required.

5 Conclusion

Successful operation of the MER vehicles has depended on both manually-directed and autonomous driving. The two methods are complementary, and careful selection of the right techniques leads to better overall performance in the face of limited time, power, imagery, and onboard computation.

While most of the distance covered by both rovers has been on level ground with varying degrees of geometric hazards, most of the time has been spent in more challenging environments coupling steep slopes with loose materials and positive obstacles. Careful terrain analysis is required in these cases, and VisOdom has also been absolutely essential for safe and accurate driving.

6 Acknowledgements

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology under contract to the National Aeronautics and Space Administration. Thousands of people were involved in the mission at JPL and throughout the world, in government, academia, and industry. We gratefully acknowledge their outstanding work, which enabled us to explore Mars through Spirit and Opportunity. We would like to thank the entire Mars Exploration Rover team, and in particular the other rover planners: Khaled Ali, Eric Baumgartner, Paolo Bellutta, Bob Bonitz, Brian Cooper, Frank Hartman, Scott Maxwell, Ashley Stroupe, Ashitey Trebi-Ollennu, Eddit Tunstel, John Wright, and Jeng Yen.

References

[1] Jeffrey Biesiadecki, Mark Maimone, and Jack Morrison. The Athena SDM rover: A testbed for Mars rover mobility. In *International Symposium on Artificial Intelligence, Robotics, and*

Automation for Space (i-SAIRAS), Montreal, Canada, June 2001. http://robotics.jpl.nasa.gov/people/mwm/sdm-mobility/.

[2] Yang Cheng, Mark Maimone, and Larry Matthies. Visual odometry on the Mars Exploration Rovers. In *IEEE Conference on Systems, Man and Cybernetics*, The Big Island, Hawaii, USA, October 2005.

[3] Paul Backes et al. Sequence planning for the Fido Mars rover prototype. In *IEEE Aerospace Conference*, Big Sky, Montana, USA, 2003.

[4] Steven B. Goldberg, Mark W. Maimone, and Larry Matthies. Stereo vision and rover navigation software for planetary exploration. In *IEEE Aerospace Conference*, volume 5, pages 2025–2036, Big Sky, Montana, USA, March 2002. http://robotics.jpl.nasa.gov/people/mwm/visnavsw/.

[5] M. Golombek and D. Rapp. Size-frequency distributions of rocks on Mars and Earth analog sites: Implications for future landed missions. *Journal of Geophysical Research - Planets*, 102(E2):4117–4129, February 1997.

[6] C. Leger and R. Deen. Remote image analysis for Mars Exploration Rover mobility and manipulation operations. In *IEEE Conf. on Systems, Man and Cybernetics*, Big Island, Hawaii, USA, October 2005.

[7] Rongxing Li, Brent A. Archinal, Raymond E. Arvidson, Jim Bell, Philip Christensen, Larry Crumpler, David J. Des Marais, Kaichang Di, Tom Duxbury, Matt Golombek, John Grant, Ronald Greeley, Joe Guinn, Andrew Johnson, Randolph L. Kirk, Mark Maimone, Larry H. Matthies, Mike Malin, Tim Parker, Mike Sims, Shane Thompson, Steven W. Squyres, and Larry A. Soderblom. Spirit rover localization and topographic mapping at the landing site of Gusev Crater, Mars. *JGR-Planets, Special Issue on Spirit Rover,* to appear, 2005.

[8] J. N. Maki, J. F. Bell III, K. E. Herkenhoff, S. W. Squyres, A. Kiely, M. Klimesh, M. Schwochert, T. Litwin, R. Willson, A. Johnson, M. Maimone, E. Baumgartner, A. Collins, M. Wadsworth, S. T. Elliot, A. Dingizian, D. Brown, E. C. Hagerott, L. Scherr, R. Deen, D. Alexander, and J. Lorre. Mars exploration rover engineering cameras. *Journal of Geophysical Research*, 108(E12):12–1–12–24, December 2003. http://www.agu.org/pubs/crossref/2003/2003JE002077.shtml.

[9] Reid Simmons, Lars Henriksen, Lonnie Chrisman, and Greg Whelan. Obstacle avoidance and safeguarding for a lunar rover. In *AIAA Forum on Advanced Developments in Space robotics*, Madison, WI, August 1996. http://www.cs.cmu.edu/˜reids/papers/AIAAobsAvoid.pdf.

[10] J. Yen, B. Cooper, F. Hartman, S. Maxwell, and J. Wright. Sequence rehearsal and validation on surface operations of the Mars Exploration Rovers. In *SpaceOps*, Montreal, Canada, 2004.